



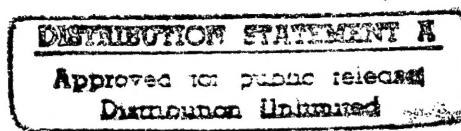
Department of Computer
and Information Science

2015 Neil Avenue
Columbus, OH 43210-1277

Phone 614-292-1553
FAX 614-292-2911
Email fz@cis.ohio-state.edu

October 1, 1997

Dr. Michael F. Shlesinger
ONR 331
Office of Naval Research
Ballston Center Rower One
800 North Quincy Street
Arlington, VA 22217-5660



Dear Mike,

I would like to update you with our latest research accomplishments in the project funded under the ONR YI Award (ONR Grant No. N00014-97-1-0599). Our project, titled [Practical control algorithms for nonlinear dynamical systems using phase-space knowledge and mixed numeric and geometric computation], aims to develop novel computational tools for synthesizing and analyzing controllers for a class of complex physical systems.

During the period June 1 — September 30, 1997, we have obtained the following research results:

- Developed empirical performance criteria for characterizing stabilities and robustness of the maglev control experimental system and completed a preliminary implementation of the performance characterization algorithms. A graduate student, Shiou Loh, has just completed a Masters' thesis on this subject.
- Developed and experimented with a phase-space search algorithm for synthesizing control actions. Another graduate student, Jeff May, has been working with me on this problem and is in the process of implementing additional performance metrics for measuring control performance.
- Presented part of the above results in an article, "Phase-Space Nonlinear Control Toolbox: The Maglev Experience" at *HS'97: Fifth International Hybrid Systems Workshop*, Notre Dame, IN, Sept. 11-13, 1997. We will also report our work at the *AAAI Fall 1997 Symposium on Model-Directed Autonomous Systems* in Boston, MA, Nov 7-9, 1997.

I have enclosed a copy of the above mentioned article. Please do not hesitate to contact me if you need additional information.

Best Regards,

A handwritten signature in black ink, appearing to read "Feng Zhao".

Feng Zhao
Associate Professor

19971014 061

DIGITAL QUALITY INSPECTED 8

Phase-Space Nonlinear Control Toolbox: The Maglev Experience

Feng Zhao

Department of Computer and Information Science
The Ohio State University
Columbus, OH 43210
`{fz,may-j}@cis.ohio-state.edu`

Shiou C. Loh

Jeff A. May

Abstract

We describe the Phase-Space Nonlinear Control Toolbox, a suite of computational tools for synthesizing and evaluating control laws for a broad class of nonlinear dynamical systems. The Toolbox comprises computational algorithms for identifying optimal control reference trajectories in the phase space of dynamical systems and experimental methods for evaluating performance of the control laws. These algorithms combine knowledge of the geometric theory of modern nonlinear dynamical systems with efficient computational methods for geometric reasoning and graph search; they define the properties of controllability and robustness in terms of phase-space geometric structures and exploit the phase-space neighborhood adjacencies to obtain computational efficiency. Compared to the traditional analytic control design methods, the phase-space based control synthesis and evaluation rely on high-performance computational techniques and are applicable to physical systems operating in large nonlinear regimes. Using a proof-of-concept physical experiment for stabilizing a nonlinear magnetic levitation system, we have successfully demonstrated the feasibility of the phase-space control technology.

Introduction

Many physical systems such as man-made electro-mechanical systems operate in large nonlinear regimes. These systems often exhibit extremely complex behaviors that defy conventional analytical analysis and numerical simulations.

This paper describes a suite of computational methods for nonlinear control synthesis and analysis that computationally explore the phase space of dynamical systems guided by domain knowledge of dynamical systems and control theory. We have constructed a proof-of-concept physical experiment for a nonlinear magnetic levitation (maglev) control system and demonstrated the feasibility of the phase-space control technology. The maglev project serves as a testbed for developing practical phase-space based control al-

gorithms, performance criteria, and evaluation methods.

The phase-space based control has several important advantages compared to conventional control design methods. It relies on a computational characterization of phase-space geometry and exploits the geometric knowledge to guide control planning and execution. It captures global dynamical behaviors and requires no linear approximation. It explores a much larger space of possible control strategies than linear control methods do. It trades computational resources with control performance. The important stability and robustness properties of a control system are geometrically interpreted in phase space and can be operationally verified using the geometric models. We expect the phase-space control methods to complement conventional techniques and find niches where conventional methods are not applicable.

Phase-Space Control Synthesis

Poincare's geometric method of modern dynamical systems provides the theoretical basis for the phase-space analysis and synthesis of nonlinear dynamics (Guckenheimer & Holmes 1983). A phase space for a dynamical system is spanned by the independent state variables of the system. For instance, a swinging planar pendulum's phase space is a two-dimensional plane of position versus velocity. In phase space, important qualitative behaviors of dynamical systems are characterized by the geometric features of the space such as points, curves, surfaces, and volumes — equilibrium points, limit cycles, stability regions, trajectory flows, and their spatial arrangement — that can be extracted, identified, and exploited through computational means.

Control theory and engineering provide a body of tools for designing linear control systems. Mathematical results on stability and controllability of linear systems have successfully guided practical implementations of linear controllers. In contrast, nonlinear control lacks general methods that provide a unified treatment of and approach to a wide class of nonlinear systems. Recent theoretical work on the controllability of

nonlinear systems that employs a differential geometric approach is still far from being practical (Isidori 1985). Conventional analytical methods have two important limitations: they require accurate models and labor-intensive simulation and calibration.

- As an example of a nonlinear control system, we examine the well-studied control problem for legged robots. The hopping motion of a legged robot can be described as a limit cycle in a phase space of position and velocity, with distinct points (events) delineating phases such as lift-off, touch-down, or flight phase of the motion (Raibert 1986). The behaviors of the robot are characterized by the number and shapes of the limit cycle. After having successfully designed several generations of legged robots, Marc Raibert anticipated the need for and the possibility of automatic control synthesis methods for designing more complex legged machines such as gymnastic robots:

The strategy we chose is based on several decisions... Each of these decisions was made by humans based on knowledge of the mechanics of the problem and intuition. It is not hard to imagine that future control systems may be able to formulate strategies such as these automatically (Hodgins & Raibert 1987).

The phase-space approach provides a systematic way to explore the control spaces of physical systems. Radically different control behaviors can be automatically synthesized and evaluated by systematically varying operating conditions and initial states of the robots. For instance, by varying the magnitude and length of the thrust for the robot, distinct trajectories may be generated, examined, and selected according to specified optimality criteria.

Overview of phase-space control

Combining the phase-space geometric description of dynamics with mathematical characterization of stability and controllability from control theory, we have developed computational algorithms for aggregating, classifying, and searching for optimal control reference trajectories (Bradley & Zhao 1993; Zhao 1991; 1994). The main ideas of the phase-space control synthesis are illustrated here using the spatial aggregation framework (Yip & Zhao 1996). For simplicity, we consider the stabilization control problem where the control objective is to find reference trajectories to steer the system towards a prespecified goal state in phase space. Given the phase-space data descriptions as input, the spatial aggregation operators **aggregate**, **classify**, and **search** transform the low-level field data into global control policies by exploiting the spatial-temporal neighborhood structures of the vector fields.

- Aggregation:** A phase-space vector field, parameterized for a particular control action, are aggregated

into a neighborhood graph explicitly encoding adjacencies for a grid of phase-space cellular regions. At the higher level, neighborhood graphs parameterized for different control actions are aggregated to a composite neighborhood graph. Two cells are adjacent if there is a trajectory that connects the cells under a certain control action.

The nodes of neighborhood graph can further record probabilities of transition due to discretization or uncertainties in measurements. It is possible to have multiple edges directly connecting two cells due to different control actions. The edges of the graph can be weighed according to the quality of control. For instance, the weight can measure the amount of control resources such as time or control energy consumed while making the transition.

- Classification:** A neighborhood graph is classified into equivalence classes of cells. At the first level, the cells are grouped into behavioral classes according to robustness or other control concerns. Information such as the likelihood of being perturbed off the path is valuable for the search procedure at the higher level.

At the next level, the cells in the composite neighborhood graph are classified into two classes: one corresponds to the collection of cells in the controllable region, i.e., each cell has a path to the goal state; the other class comprises cells that are disconnected from the goal. A more sophisticated classification scheme labels groups of cells according to the probability of making it to the goal starting at different initial states in a cell. We have implemented a version that classifies regions into controllable, marginally controllable, and uncontrollable types.

- Search:** The classified phase-space neighborhood structure is then searched for optimal control trajectories. A control reference trajectory (or a control law) consists of a sequence of cells and the associated control actions that induce transitions from cell to cell.

We have implemented an iterative algorithm that computes optimal paths based on a given optimality function (e.g. based on response time or robustness). Additionally, we have implemented a dynamic programming algorithm that efficiently identifies shortest paths to the goal in the graph according to various definitions of distance functions.

The control paths can be computed using the system model or state measurements taken directly from the system. The paths are then stored in a table for efficient run-time retrieval. Sub-optimal control paths can also be retained to permit graceful control degradation in the presence of uncertainties. The controller can trade off the amount of computation with run-time control quality.

The main steps of aggregation, classification, and search are summarized in Figure 1.

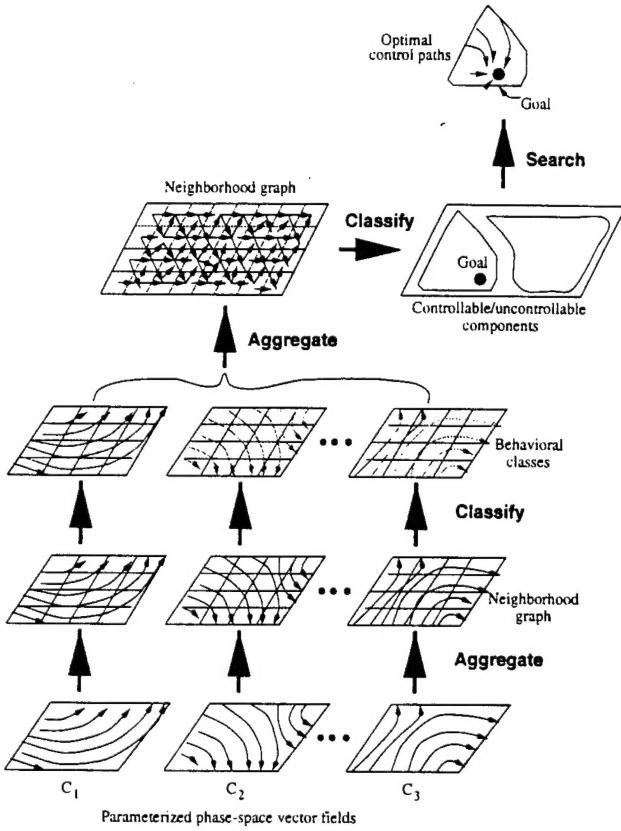


Figure 1: Spatial aggregation of phase-space control trajectories. The phase-space trajectories, parameterized for a particular control action, are first aggregated into a neighborhood graph and classified into behavioral groups according to various definitions of behavioral equivalence. Given admissible control actions, the classified neighborhood graphs are aggregated to form a composite neighborhood graph which can be then be classified and searched out to identify optimal control policies.

Geometric interpretation of control performance

We formalize the important control properties for nonlinear dynamical systems using the phase-space geometric structures. More specifically, we define controllability, stability, robustness, and optimality of control systems geometrically so that these properties can be operationally verified.

- The controllability criterion tests if a particular target state is reachable from a given initial state or an operating region of a system under a control signal. The set of reachable states form a subspace of phase space that can be computed geometrically.
- The stability criterion tests if a system remains within the neighborhood of a target state for a reasonable amount of time. Stability is characterized by the stability region—the set of initial states that evolve to the same limit set. The stability region is a subspace of phase space that can also be computed geometrically.
- The robustness criterion checks if a system attains the same properties when parametric or structural uncertainties are introduced. Geometrically, certain types of uncertainties like noise or measurement errors can be modeled as regions around states or a sequence of states.
- The optimality criterion tests if the system achieves the goal with an optimal amount of resources such as time and energy. The consideration on the resource consumption parameterizes phase-space trajectories. Control reference trajectories are ranked according to prespecified optimality metric.

The phase-space geometric definition of controllability and robustness becomes the basis for developing the experimental evaluation methods for the maglev control system.

The Maglev Control Experiment

Magnetically levitated vehicles such as the German Transrapid train (Figure 2(a)) require active suspension technology that maintains a constant air gap and dampens disturbances caused by road irregularities and wind loads (Eastham 1989). Figure 2(b) shows a maglev vehicle and guideway system with active primary and secondary suspensions (McCallum & others 1992). In such a system, the control objective for the active suspension system is to ensure a good ride quality for passengers on realistic guideway systems.

In the earlier work, we had developed a phase-space control algorithm that searches through equivalence classes of trajectories to synthesize a global switching control law (Zhao 1991; Zhao & Thornton 1992). Simulation result shows that the synthesized controller compares favorably to classical linear feedback design and does not require linear approximations to the model dynamics. The current maglev control project aims

at developing practical phase-space control algorithms and implementations for a prototype of maglev system whose dynamics resembles that of the German Transrapid. The implemented control system monitors the state of the maglev system, computes the required control action, and actuates an electromagnet to counter disturbances.

A maglev control system prototype

We have constructed a prototype for the maglev control experiment in the lab, shown in Figure 3. The system uses an electromagnet to suspend a steel ball in the air and is representative of magnetically levitated transportation systems (EMS system) such as the German Transrapid system. At the equilibrium point, the attracting magnetic force balances the gravitational force acting on the ball. However, this attractive system is inherently unstable¹. An active controller is required to maintain a constant air gap between the ball and the magnet in the presence of disturbances.

The block diagram of the experimental control system is shown in Figure 4. The data for ball displacement, velocity, and solenoid current is collected through photo sensors and current sensor. A 12-bit analog-to-digital converter samples the data at a rate of about 5000Hz. A digital computer (Pentium 75MHz) implements the control algorithm and a low-pass filter, and provides a run-time user interface. The digital controller employs either a global or a local control algorithm, depending on the current state in phase space. The appropriate control signal is delivered to the digital-to-analog converter, amplified, and then applied to the electro-magnet that suspends the ball. While the control is in progress, the user can interrupt the system through the user interface for tasks such as introducing disturbance.

The maglev model

The *nonlinear* model for the maglev system is described by

$$\begin{cases} \frac{dx}{dt} = v \\ \frac{dv}{dt} = g - \frac{L_0 x_0 I^2}{2 m x^2} \end{cases} \quad (1)$$

where the state variables x and v represent the vertical gap between steel ball and magnet and the vertical velocity of the ball, respectively. The control parameter is the coil current I . The other parameters are the ball mass $m = 0.008432Kg$, the solenoid-ball system inductance $L_0 = 0.00802H$ at the equilibrium, the desired equilibrium vertical gap $x_0 = 0.0066m$, and the gravitational acceleration $g = 9.81m/s^2$. The nonlinearity of the system comes from the inverse square magnetic force law.

¹In contrast, an EDS maglev system uses repulsive magnetic force to suspend vehicles and is inherently stable. However, an EDS system requires superconducting circuits in order to reduce energy loss.

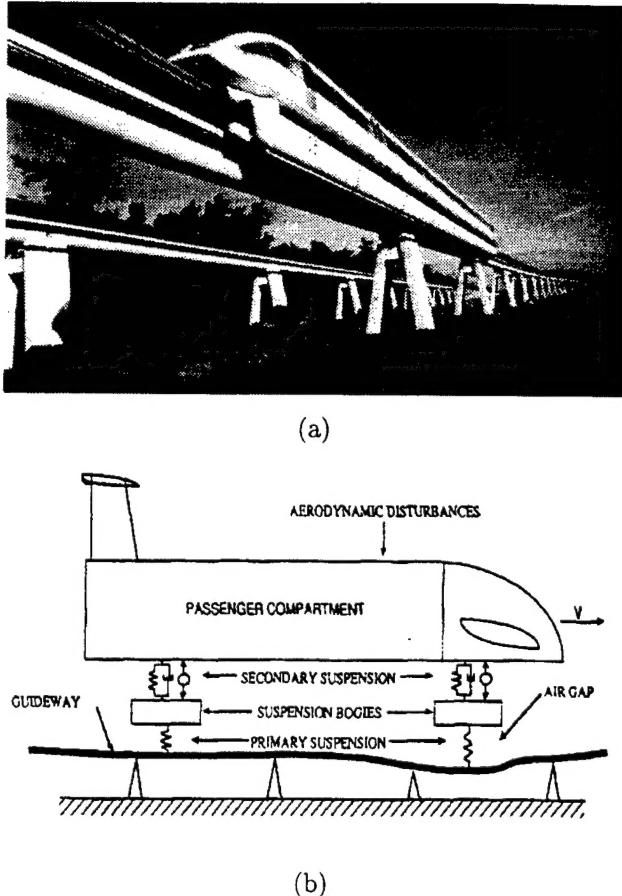


Figure 2: (a) German Transrapid magnetically levitated transportation system. (b) A maglev vehicle and guideway system with active primary and secondary suspensions, subject to aerodynamic disturbances and guideway variations. Such a system, often highly nonlinear, requires a high-performance controller to ensure ride smoothness.

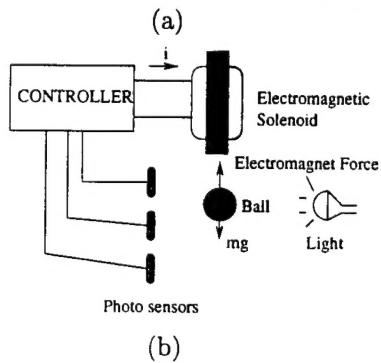
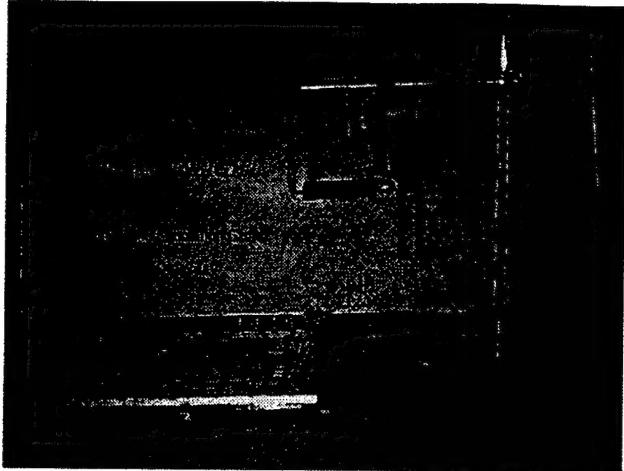


Figure 3: The maglev control system prototype: (a) Photo of the physical experiment. (b) Basic components.

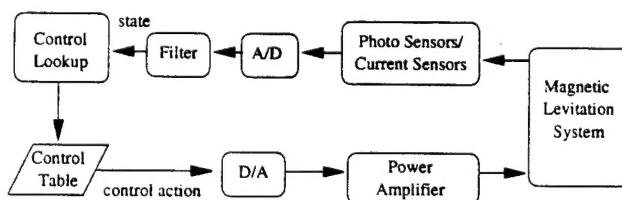


Figure 4: Block diagram for the maglev control system.

At the desired equilibrium gap x_0 , there is a unique coil current I_0 for which the magnetic force exactly counterbalances the force due to gravity and the ball has no vertical velocity and acceleration. However, the equilibrium is a saddle node which is not stable. The control objective, therefore, is to stabilize the ball and maintain a constant air gap despite any disturbances. The available control input is the coil current I in the model (1) provided by a voltage-controlled current source, varying from 0.03A to 0.83A. The power supply delivers $I_0 = 0.38A$ at the equilibrium.

Phase-Space Control Trajectory Design: an algorithm

The phase-space based control algorithm determines a control law for a discrete sampled system. The algorithm assumes that the control objective is to move the system to a prespecified goal state (or set of states) and that the behavior of the system is governed by the following map:

$$x_{n+1} = F(x_n, c)$$

where x_n is the state vector of the system at a given time, c is the control input vector, and x_{n+1} is the state of the system after one sampling period (it is assumed that the sampling period is fixed). If the behavior of the system is described in continuous terms, the discrete equation can be obtained by integrating the continuous equation over one sampling period.

The algorithm takes as input the model for the system, the goal state, the set of allowable control values, and information about how to partition the phase space. It outputs a table of control actions, with each cell having its own entry in the table. The runtime control program determines which cell contains the current system state and indexes into the control table to retrieve the appropriate control output.

The algorithm consists of three main steps.

1. Partition the phase space into cellular regions².
2. Compute the graph of cell adjacencies. A cell x is said to be adjacent to a cell y if, for some allowable control value c , there exists a natural number n such that $F^k(\text{center}(x), c) \in x, \forall k < n$ and $F^n(\text{center}(x), c) \in y$.
3. For each cell c , find a path to the goal cell (i.e. the cell containing the goal state). If no such path exists, c is marked as outside the controllable region of the system. Otherwise, choose one such path, and enter the control output corresponding to this path into the control table.

Step three employs path selection algorithms to choose a path to the goal cell when multiple such paths exist. Typically, path selection will be based on system performance criteria. For example, if a fast system response is desired, short paths can be chosen over longer paths.

²The phase-space partition does not have to be uniform.

Phase-Space Control Performance Evaluation

Most nonlinear systems are not amenable to analytic characterization of control performance such as controllability and robustness. The phase-space control method is particularly well suited for synthesizing control for this class of systems. The phase-space interpretation of control properties provides a basis for computational implementation of practical control performance evaluation methods for nonlinear systems. These evaluation methods can be used to experimentally validate if an implementation satisfies the design specification and compare different control strategies.

- Controllable region and robustness

The controllable region is defined as the collection of cells from which the goal states are reachable. Operationally, the region is characterized by aggregating cell transitions using measured trajectories. The controllable region can be labeled with sizes and other more refined performance characterization.

Robustness measures the ability of the system to withstand disturbances. Experimentally, we characterize for each state of the system the magnitude and duration of destabilizing disturbance. In the experiment, the disturbance is introduced as part of the control input to the system.

A system with a larger controllable region is generally more stable and robust to disturbances such as a sudden displacement.

- Settling time, rise time, and overshoot

Settling time is defined as the time for the system to settle in the vicinity of the steady state. That is, it measures the time required for the transient state to decay. On the other hand, rise time is the time for the system output to reach the vicinity of the desired point for the first time, and measures the response time of the system. Overshoot is the maximum amount the system output overshoots the final value. The amount of overshoot is a good indication of the smoothness of control and is normally measured as the percent relative to the final value. A control design typically requires trade-off between response time and overshooting.

Experimentally, we define a set of iso-curves for settling time with respect to a steady state in phase space, shown in Figure 5. Rise time is likewise characterized. Overshoot is indirectly characterized by the curvature along trajectories in the neighborhood of an equilibrium.

Experimental Results and Analysis

We have implemented and evaluated the phase-space control algorithm on the maglev system. In the experiment, the phase space of the maglev system is partitioned into a 50 by 50 uniform cellular space. The

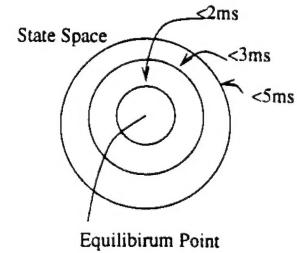


Figure 5: Iso-curves for settling time with respect to a stable equilibrium state, shown in phase space.

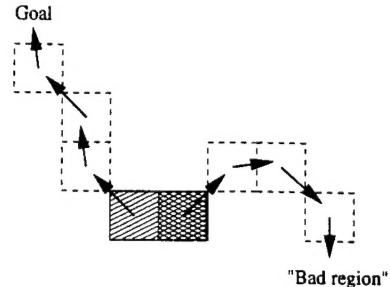


Figure 6: Two adjacent phase-space cells at a behavioral boundary. Each arrow, for a given control value, indicates the direction of the field at the cell center. The two shaded cells lie on a behavioral boundary since their field directions differ substantially.

control space is also sampled to form a discrete control set. The path search algorithm uses an iterative order cn^2 method (where n is the number of cells in the phase space discretization and c is the number of divisions in the control space) that does not require computation of all possible paths. Paths are selected on the basis of two criteria. The first criteria gives preference to shorter paths (with respect to the number of cells contained in the path) over longer paths. The second criteria gives preference to paths that are more robust with respect to disturbances or modeling errors. More specifically, it prefers paths that are further away from the "behavioral boundaries" of the system. For our purposes, a behavioral boundary in phase space is an area of phase space where the field direction in neighboring cells differs substantially (Figure 6). Paths that use cells close to a behavioral boundary are more susceptible to disturbances or modeling error, since a small error can result in a large change in the expected trajectory near these boundaries.

Figure 7 shows the control graph generated for the magnetic levitation system. The control table is computed off-line. Real-time control is accomplished by fast sensing, state estimation, and control action lookup. The phase-space control incorporates a local control that takes over when the system enters the neighborhood of the equilibrium.

The experimental results on the physical maglev sys-

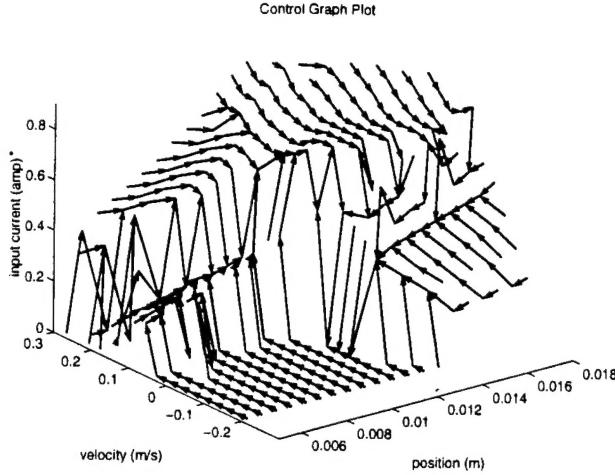


Figure 7: Synthesized control paths (to the goal) are plotted against phase space, with lines indicating paths traveled for each cell. The phase space is divided into a uniform 20 by 20 grid for the purpose of visualization. Continuous lines indicate the chosen path from a cell to the goal state. The height of a point along a path corresponds to the magnitude of the control input at that location.

tem show that the global, phase-space control algorithm is able to suspend the steel ball reliably under normal operating conditions for at least 15 minutes (after 15 minutes the apparatus must be turned off to avoid overheating of the solenoid). We have also introduced external disturbances into the system (in the form of a momentary current drop that results in a sudden change in air gap and velocity) and experimentally measured the robustness of the system.

The phase space plot of the controlled trajectory is shown in Figure 8. Figure 9 shows how the air gap, velocity, and control vary over time. The initial bump in the air gap plot is due to the disturbance artificially introduced. The controller is activated to stabilize the system as soon as the disturbance is detected.

Using the phase-space performance evaluation method, the controllable region for the global controller is depicted in Figure 10.

Related work

The field of nonlinear control has developed a rich collection of design methods that apply to specific classes of problems (Slotine & Li 1991). For instance, methods of linearization and describing functions generalize linear techniques to nonlinear systems. The technique of gain scheduling, for example, approximates a nonlinear control system with a piecewise-linear one and designs a linear controller for each linear piece. To apply the gain scheduling technique, one has to decompose the phase space into locally linear regions and design a linear controller for each region. Phase-space con-

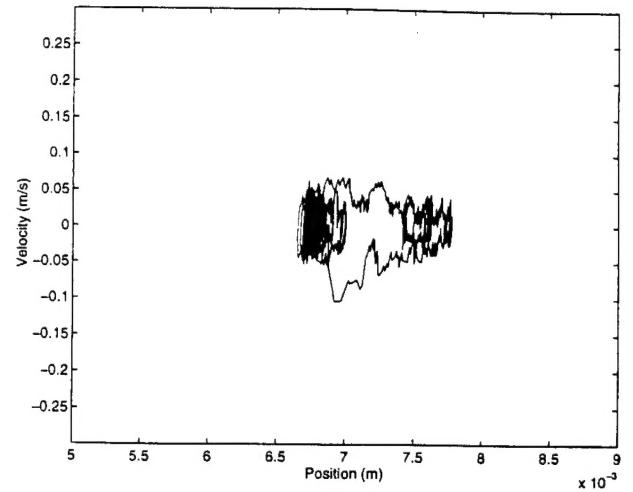


Figure 8: A sample maglev control trajectory shown in phase space (position \times velocity). The equilibrium point $(0.0066, 0)$ is inside the dense region covered by the orbit. Notice that the global controller is able to bring the trajectory back to equilibrium after an initial disturbance.

trol generalizes the decomposition process to permit a systematic computational exploration.

Hsu developed a cell-based method called cell-to-cell mapping for approximating state spaces (Hsu 1987). A continuous state space is discretized into regular cells forming a cell space. The associated map of a system becomes a cell-to-cell map which maps one cell to another cell. The cell-to-cell mapping method approximates the stability region of an attracting cell with a collection of cells that eventually map to that cell. Our phase-space control algorithm goes a major step beyond the cell-based analysis. We had previously developed a phase-space control framework for exploiting phase-space global knowledge to obtain high-quality control design (Bradley & Zhao 1993). This work focuses on developing practical methods and implementations for the phase-space control technology. We use the cell decomposition as a first order approximation to the phase-space geometry to develop practical control synthesis and evaluation methods. The cellular neighborhood structure serves as a place holder for phase-space control performance data such as resource consumption, controllability, and robustness that can be actively exploited in synthesizing high-quality control actions. Additionally, the structure permits programmed trade-offs between computation and control quality.

Research in hybrid systems, a class of dynamical systems that possess both continuous dynamics and discrete transitions, has produced a body of theories and algorithms for control analysis and synthesis (Asarin, Maler, & Pnueli 1995; Caines & Wei 1995; Branicky 1995; Brockett 1993; Henzinger *et al.* 1995;

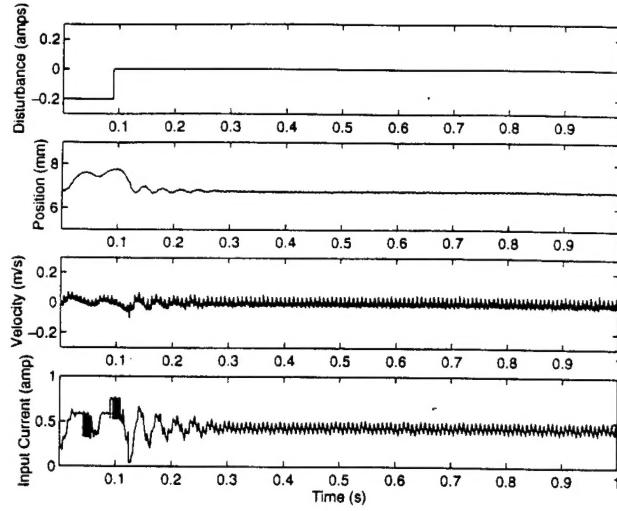


Figure 9: Experimental data from the maglev system: (a) Disturbance introduced at the beginning. (b) Air gap x vs. time. Notice that the system stabilizes after the initial disturbance. (c) Velocity v vs. time. (d) Control I vs. time. The global, phase-space controller is activated whenever the position or velocity of the ball deviates from the desired value significantly. When the system is within the capture region of the local control, the controller switches to the local control.

Nerode & Kohn 1993; Stiver, Antsaklis, & Lemmon 1995). Phase-space control systems form a special class of hybrid systems. The formal results obtained in hybrid systems research provide theoretical characterizations for the phase-space control systems. We have focused on developing practical algorithms for synthesizing control actions for nonlinear systems, using phase space as a geometric model. We have also characterized the computational complexity for a restricted class of hybrid systems (Kolen & Zhao 1997).

Recently, neural net technology has been successfully applied to the control of highly nonlinear systems. In one case, a neural net is used to model the inverse dynamics of a high-performance fighter jet in order to produce high-quality control design (Totah 1996). However, one major drawback of the neural-net based controller is its lack of stability and performance guarantee required by the FAA certification process. The tools described in this paper can computationally characterize operational properties of these nonlinear controllers that are not amenable to traditional analytical analysis. Hence, our approach is complementary to existing nonlinear control technology.

Conclusion

We have described the Phase-Space Nonlinear Control Toolbox for synthesizing and evaluating control laws in phase space. We have demonstrated, using the maglev

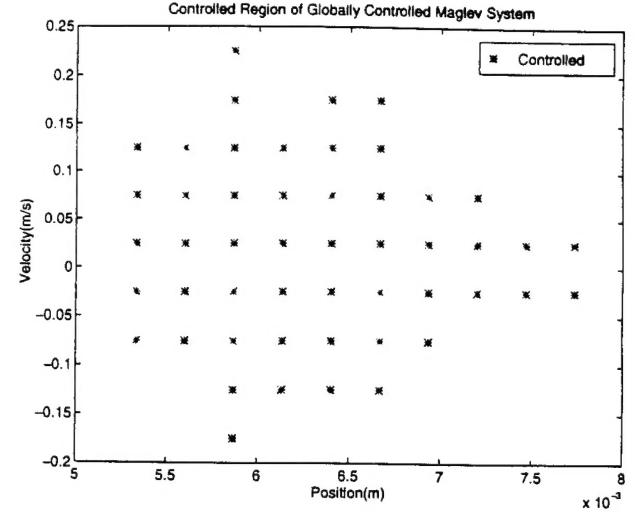


Figure 10: Region of controllability for the phase-space based global control law, comprising cells marked by *.

experiment, that the phase-space control technology is feasible. The current capability of the Toolbox is limited in a number of ways. We plan to incorporate and exploit additional global phase-space features, new optimization metrics, statistics-based evaluation methods, and learning into the phase-space framework. The phase-space control is applicable to a broad class of physical systems operating in large nonlinear regimes, for which conventional linear, analytic methods are ill-suited.

The rapid advances in information processing, micro-fabrication, and MEMS are fueling a new generation of distributed autonomous systems (Williams & Nayak 1996). Our abilities to sense and act in the complex physical environments are increasingly augmented by massive networks of tiny, invisible sensors, actuators, and computers embedded in everything from appliances to materials to building structures. These distributed computational agents are immersed in physical media and governed by the fundamental laws of computation and physics. Computational methods such as phase-space control may enable these agents to maximumly exploit the environment at the juncture of digital universe and continuous physical world.

Acknowledgments

This work is supported in part by ONR Young Investigator Award N00014-97-1-0599, NSF Young Investigator Award CCR-9457802, NSF grant CCR-9308639, Alfred P. Sloan Foundation Research Fellowship, and a grant from Xerox Palo Alto Research Center.

References

Asarin, E.; Maler, O.; and Pnueli, A. 1995. Reachability analysis of dynamical systems having piecewise-

constant derivatives. *Theoretical Computer Science* 138:35–65.

Bradley, E., and Zhao, F. 1993. Phase-space control system design. *IEEE Control Systems* 13(2):39–47.

Branicky, M. 1995. Universal computation and other capabilities of hybrid and continuous dynamical systems. *Theoretical Computer Science* 138:67–100.

Brockett, R. 1993. Hybrid models for motion control systems. *Essays in Control: Perspectives in the Theory and its Applications*.

Caines, P. E., and Wei, Y.-J. 1995. On dynamically consistent hybrid systems. *Hybrid Systems II, Lecture Notes in Computer Science*.

Eastham, T. R. 1989. Maglev: A realistic option for the nineties. *Magnetic Levitation Technology for Advanced Transit Systems*.

Guckenheimer, J., and Holmes, P. 1983. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, New York.

Henzinger, T.; Kopke, P.; Puri, A.; and Varaiya, P. 1995. What's decidable about hybrid automata? In *Proceedings of the 27th Annual Symposium on Theory of Computing*, 373–382. ACM Press.

Hodgins, J., and Raibert, M. 1987. Biped gymnastics. *Fourth International Symposium of Robotics Research*.

Hsu, C. S. 1987. *Cell-to-Cell Mapping*. Springer-Verlag, New York.

Isidori, A. 1985. *Nonlinear Control: An Introduction*. Springer-Verlag, New York.

Kolen, J., and Zhao, F. 1997. A computational analysis of the reachability problem for a class of hybrid dynamical systems. *Hybrid Systems IV, Lecture Notes In Computer Science*.

McCallum, D., et al. 1992. Aerodynamic control of maglev vehicle. *A talk given at the MIT Maglev Consortium Symposium*.

Nerode, A., and Kohn, W. 1993. Models for hybrid systems: automata, topologies, controllability, observability. *Hybrid Systems, Lecture Notes In Computer Science*.

Raibert, M. 1986. *Legged Robots That Balance*. MIT Press.

Slotine, J. E., and Li, W. 1991. *Applied Nonlinear Control*. Prentice-Hall, Englewood Cliffs, New Jersey.

Stiver, J.; Antsaklis, P.; and Lemmon, M. 1995. An invariant based approach to the design of hybrid control systems containing clocks. *Hybrid Systems III, Lecture Notes in Computer Science* 1066:464–474.

Totah, J. 1996. Simulation evaluation of a neural-based flight controller. In *Proc. of AIAA*.

Williams, B., and Nayak, P. 1996. Immobile robots: AI in the new millennium. *AI Magazine*.

Yip, K., and Zhao, F. 1996. Spatial aggregation: Theory and applications. *Journal of Artificial Intelligence Research* 5:1–26.

Zhao, F., and Thornton, R. 1992. Automatic design of a maglev controller in state space. In *Proc. of the 31st IEEE Conf. on Decision and Control*.

Zhao, F. 1991. Phase space navigator: Towards automating control synthesis in phase spaces for nonlinear control systems. In *Proc. 3rd IFAC Int'l Workshop on AI in Real Time Control*. Pergamon Press.

Zhao, F. 1994. Extracting and representing qualitative behaviors of complex systems in phase spaces. *Artificial Intelligence* 69(1–2):51–92.